



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/808,561	03/24/2004	Ofir Zohar	TUC920085005US1 (0131.U01)	4776
85071 7590 08/06/2009 GRIFFITHS & SEATON PLLC (IBM2) 2108 N. Lemon Street Mesa, AZ 85215				
EXAMINER BRADLEY, MATTHEW A				
ART UNIT		PAPER NUMBER		
2187				
NOTIFICATION DATE		DELIVERY MODE		
08/06/2009		ELECTRONIC		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

docketing@gs-iplaw.com

Office Action Summary

Application No.

10/808,561

Applicant(s)

ZOHAR ET AL.

Examiner

MATTHEW BRADLEY

Art Unit

2187

Period for Reply -- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 04 May 2009.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 41-73 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 41-73 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-946)
- 3) ☐ Information Disclosure Statement(s) (PTO/SG/08)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

Response to Amendment

This Office Action has been issued in response to amendment filed 4 May 2009. Applicant's arguments have been carefully and fully considered and are partially persuasive. Therefore, the rejection has been withdrawn. Accordingly, this action has **NOT** been made final.

Claim Status

Claims 41-73 remain pending and are ready for examination.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims **41-48** and **52-61**, are rejected under 35 U.S.C. 102(e) as being anticipated by Hicken et al. (US 2004/0153727).

As per independent claim **41**, Hicken et al. disclose a method for managing a data storage system (300 of Fig. 3; paragraph 0038, lines 4-10), comprising:

- configuring a first cache (339 of Fig. 3) to perform at least one of the operations of retrieving data from and storing data at a first range of logical addresses (LAs) in a storage device (paragraph 0038, lines 13-17; paragraph 0039, lines 7-10;

paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2);

- configuring a second cache (333 of Fig. 3) to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (paragraph 0039, lines 5-7; paragraph 0041, lines 10-14);
- configuring one or more third caches (338 of Fig. 3) to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device the first, second, and third caches being implemented in three separate physical units comprising software and adapted to function as controllers substantially independently of each other; (paragraph 0041, lines 14-17, the Examiner notes that as the caches store data, the caches store software as instantly claimed);
- detecting an inability of the second cache to retrieve data from or store data at the first range of LAs (paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails); and
- reconfiguring at least one the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (paragraph 0044).

As per dependent claim **42**, Hicken et al. disclose the method according to claim 41, and comprising configuring one or more interfaces (CPUs 331 and 336 of Fig. 3) to receive input/output (IO) requests (paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16) from host processors (310 of Fig. 3; paragraph 0038, lines 4-7) directed to specified LAs (paragraph 0025, lines 3-6) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (paragraph 0039, lines 15-16).

As per dependent claim **43**, Hicken et al. disclose the method according to claim 42, wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (paragraph 0038, lines 13-17), and wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping (paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches), and wherein detecting the inability comprises generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (paragraph 0042, lines 9-13), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (paragraph 0042, lines 9-13).

As per dependent claim **44**, Hicken et al. disclose the method according to claim 41, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises processing data in the first cache and the one or more third caches so as to restore the first cache and the one or more third caches to a state of full data redundancy (paragraph 0044, lines 15-24).

As per dependent claim **45**, Hicken et al. disclose the method according to claim 44, wherein processing the data comprises classifying data in the first cache into a plurality of data groups (paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified).

As per dependent claim **46**, Hicken et al. disclose the method according to claim 45, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the one or more third caches (paragraph 0043, lines 10-15; dirty data is stored on the third cache 338).

As per dependent claim **47**, Hicken et al. disclose the method according to claim 45, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the storage device (paragraph 0043, lines 10-15; dirty data is flushed to the storage units).

As per dependent claim **48**, Hicken et al. disclose the method according to claim 41, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (paragraph 0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1).

As per dependent claim **52**, Hicken et al. disclose the method according to claim 41, and comprising providing a system manager (host computer 310 and CPUs 331 and 336 of Fig. 3) which is adapted to configure the first, second and one or more third caches (paragraph 0025, lines 4-6; paragraph 0039, lines 15-16), to detect the inability (paragraph 0028, lines 1-4), and to reconfigure the at least one of the first cache and the one or more third caches (paragraph 0042, lines 9-13, lines 18-22).

As per dependent claim **53**, Hicken et al. disclose the method according to claim 52, wherein providing the system manager comprises incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (host computer 310 and CPUs 331 and 336 of Fig. 3), and operating the one or more manager processing units in a cooperative manner (paragraph 0040; all of the CPUs are connected and work together).

As per independent claim **54**, Hicken et al. disclose a data storage system, comprising:

- a storage device (300 of Fig. 3; paragraph 0038, lines 4-10) wherein data is stored at logical addresses (LAs);
- a first cache (339 of Fig. 3) which is configured to perform at least one of the operations of retrieving data from and storing data at a first range of LAs in the storage device (paragraph 0038, lines 13-17; paragraph 0039, lines 7-10; paragraph 0041, lines 10-14; primary cache 333 of storage controller 370-1 is the

cache for LA1, and secondary cache 339 is the redundant cache for LA1, and resides on storage controller 370-2);

- a second cache (333 of Fig. 3) which is configured to perform at least one of the operations of retrieving data from and storing data at the first range of LAs (paragraph 0039, lines 5-7; paragraph 0041, lines 10-14);
- one or more third caches (338 of Fig. 3) which are configured to perform at least one of the operations of retrieving data from and storing data at a second range of LAs in the storage device the first, second, and third caches being implemented in three separate physical units, comprising software, and adapted to function as controllers substantially independently of each other; (paragraph 0041, lines 14-17, the Examiner notes that as the caches store data, the caches store software as instantly claimed); and
- a system manager (host computer 310 and CPUs 331 and 336 of Fig. 3)
 - which is adapted to detect an inability of the second cache to retrieve data from or store data at the first range of LAs (paragraph 0042, lines 7-9; when the storage controller 370-1 fails, cache memory 339 fails as well), and
 - which reconfigures at least one the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs (paragraph 0044).

As per dependent claim 55, Hicken et al. disclose the storage system according to claim 54, and comprising one or more interfaces (CPUs 331 and 336 of Fig. 3) which are configured to receive input/output (IO) requests (paragraph 0025, lines 4-6 [a distinct but almost identical embodiment]; paragraph 0039, lines 15-16) from host processors (310 of Fig. 3; paragraph 0038, lines 4-7) directed to specified LAs (paragraph 0025, lines 3-6) and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs (paragraph 0039, lines 15-16).

As per dependent claim 56, Hicken et al. disclose the storage system according to claim 55, wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs (paragraph 0038, lines 13-17), and wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping (paragraph 0025, lines 4-6; paragraph 0026, lines 1-2; the CPUs onboard the storage controllers receive storage requests from the host, and issue them to the caches), and wherein detecting the inability comprises the system manager generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs (paragraph 0042, lines 9-13), and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping (paragraph 0042, lines 9-13).

As per dependent claim **57**, Hicken et al. disclose the storage system according to claim 54, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the first cache processing data therein and the one or more third caches processing data therein so as to restore the first cache and the one or more third caches to a state of full data redundancy (paragraph 0044, lines 15-24).

As per dependent claim **58**, Hicken et al. disclose the storage system according to claim 57, wherein processing the data comprises classifying data in the first cache into a plurality of data groups (paragraph 0043, lines 10-15; when the caches are flushed, only the dirty data is flushed, not the resident data, which shows the data was classified).

As per dependent claim **59**, Hicken et al. disclose the storage system according to claim 58, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the one or more third caches (paragraph 0043, lines 10-15; dirty data is stored on the third cache 338).

As per dependent claim **60**, Hicken et al. disclose the storage system according to claim 58, wherein one of the data groups comprises dirty data, and wherein processing the data comprises storing the dirty data at the storage device (paragraph 0043, lines 10-15; dirty data is flushed to the storage units).

As per dependent claim **61**, Hicken et al. disclose the storage system according to claim 54, wherein reconfiguring the at least one of the first cache and the one or more third caches comprises retaining an initial configuration of the first cache (paragraph

0042, lines 9-13, lines 18-22; after the second cache 333 fails, the first cache 339 retains its initial configuration and is now used to address storage requests for LA1).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims **49-51** and **62-65** are rejected under 35 U.S.C. 103(a) as being unpatentable over Hicken et al. (US 2004/0153727) in view of Karger et al. ("Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web," by in the Proceedings of the 29th ACM Symposium on Theory of Computing, Pages 654-663).

As per dependent claim **49**, Hicken et al. disclose the method according to claim 41. Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (page 5, section 4, "Consistent Hashing").

Hicken et al. and Karger et al. are analogous art because they are from the same field of endeavor, namely data caching.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the minimum redistribution in the form of consistent hashing of Karger et al. with the data caching redundancy system of Hicken et al. The motivation for doing so would have been because to prevent requiring a central server to distribute a completely updated hash table to all the machines every time a new machine is added to the network (page 2, column 2, paragraph 2 beginning with "Our second...").

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Karger et al. with Hicken et al. for the benefit of a data caching system with consistent hashing to obtain the invention as specified in claims 49-51 and 62-65.

As per dependent claim **50**, Hicken et al. disclose the method according to claim 49. Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (page 5, section 4, "Consistent Hashing").

As per dependent claim **51**, Hicken et al. disclose the method according to claim 49. Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (page 2, column 1, paragraph 5, "Our first tool, *random cache trees*...").

As per dependent claim **62**, Hicken et al. disclose the storage system according to claim 54. Hicken et al. do not disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

However, Karger et al. disclose the limitation wherein reconfiguring the at least one of the first cache and the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches (page 5, section 4, "Consistent Hashing").

As per dependent claim **63**, Hicken et al. disclose the storage system according to claim 62. Hicken et al. do not disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

However, Karger et al. disclose the limitation wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function (page 5, section 4, "Consistent Hashing").

As per dependent claim **64**, Hicken et al. disclose the storage system according to claim 62. Hicken et al. do not disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function.

However, Karger et al. disclose the limitation wherein redistribution comprises redistributing the first and the second ranges using a random number function (page 2, column 1, paragraph 5, "Our first tool, *random cache trees*...").

As per dependent claim **65**, Hicken et al. and Karger et al disclose the storage system according to claim 64 wherein the system manager comprises one or more manager processing units incorporating one or more manager processing units into at least one of the storage device, the first cache, the second cache, and the one or more third caches (host computer 310 and CPUs 331 and 336 of Fig. 3 of Hicken et al), and wherein the one or more manager processing units operate in a cooperative manner (paragraph 0040; all of the CPUs are connected and work together of Hicken et al).

Claims **66-73** are rejected under 35 U.S.C. 103(a) as being unpatentable over Hicken et al (US 2004/0153727), and in view of Henry et al. (U.S. 6,898,666).

As per independent claim **66**, Henry et al. teach a storage system, comprising:

- one or more mass storage devices, coupled to store data at respective first ranges of logical addresses (LAs) (Fig. 2; column 4, lines 10-23, lines 55-67);
- a plurality of interim fast-access-time caches (cache pools 1 and 2), each of the plurality of interim fast-access-time caches comprising software (the Examiner notes that as the caches store data, the caches store software as instantly claimed) and configured to operate as controllers substantially independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (column 5, lines 45-55); and coupled to receive data

from and provide data to the one or more -mass-storage devices having LAs within the respective second range (column 4, lines 55-59); and

- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned (column 2, lines 11-15; column 4, lines 64-67).

Henry et al does not teach a further plurality of interim fast-access-time caches to be configured.

Hicken et al teach.

- a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails (paragraph 0044).

Henry et al and Hicken et al are analogous art because they are from the same field of endeavor, namely memory storage systems.

At the time of invention, it would have been obvious to one of ordinary skill in the art, having both the teachings of Hicken et al and Henry et al to combine the redundant caches of Hicken et al into Henry et al for the benefit of restoring redundancy in the event of failure.

The Examiner notes that all of the claimed elements are known. The only difference is the combination of the known elements into one system. Implementing the

redundant cache system of Hicken et al into Henry et al would have offered the system of Henry et al the benefit of having a mirror of the data within a cache should a failure occur. Upon such failure, redundancy could be reestablished (See Hicken et al Paragraph 0010-0011). Such combination would have yielded predictable results to one of ordinary skill in the art at the time of invention thus obviating that which is instantly claimed.

Therefore it would have been obvious to combine Henry et al with Hicken et al to obtain the invention as specified in claims 66-67.

As per dependent claim **67**, Henry et al. and Hicken et al teach the storage system according to claim 66, wherein the mass storage devices comprise one or more disks (disks 1-5 of Fig. 2; column 5, lines 9-23 of Henry et al).

As per independent claim **68**, Henry et al. teach a method for storing data, comprising:

- storing the data in one or mass storage devices having respective first ranges of logical addresses (LAs) (Fig. 2; column 4, lines 10-23 & 55-67);
- assigning to each of a plurality of interim-fast-access-time caches (cache pools 1 and 2), a respective second range of the LAs, each of the plurality of interim fast-access-time caches comprising software configured to operate as controllers substantially independently of one another, (each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55, the Examiner notes that as the caches store data, the caches store software as instantly claimed);

- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (column 4, lines 55-59);
- receiving input/output (IO) requests from host processors directed to specified LAs (column 2, lines 11-15; column 4, lines 64-67); and directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (column 2, lines 11-15; column 4, lines 64-67).

Hicken et al teach,

- wherein a further interim fast-access-time cache is adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data. from and provide data to the one or more mass storage devices having. LAs within the respective second range when any interim fast-access-time cache fails (paragraph 0044).

Henry et al and Hicken et al are analogous art because they are from the same field of endeavor, namely memory storage systems.

At the time of invention, it would have been obvious to one of ordinary skill in the art, having both the teachings of Hicken et al and Henry et al to combine the redundant caches of Hicken et al into Henry et al for the benefit of restoring redundancy in the event of failure.

The Examiner notes that all of the claimed elements are known. The only difference is the combination of the known elements into one system. Implementing the redundant cache system of Hicken et al into Henry et al would have offered the system

of Henry et al the benefit of having a mirror of the data within a cache should a failure occur. Upon such failure, redundancy could be reestablished (See Hicken et al Paragraph 0010-0011). Such combination would have yielded predictable results to one of ordinary skill in the art at the time of invention thus obviating that which is instantly claimed.

Therefore it would have been obvious to combine Henry et al with Hicken et al to obtain the invention as specified in claims 68-69.

As per dependent claim **69**, Henry et al. and Hicken et al teach the method according to claim 68, wherein the mass storage devices comprise one or more disks (disks 1-5 of Fig. 2; column 5, lines 9-23 of Henry et al).

As per independent claim **70**, Henry et al. teach a system for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) (Fig. 2; column 4, lines 10-23, lines 55-67), comprising:

- a plurality of interim fast-access-time caches (cache pools 1 and 2), each of the plurality of interim fast-access-time caches comprising software (the Examiner notes that as the caches store data, the caches store software as instantly claimed) and configured to operate as controllers substantially independently of one another, each interim fast-access-time device being assigned a respective second range of the LAs (column 5, lines 45-55); and coupled to receive data from and provide data to the one or more -mass-storage devices within the respective second range (column 4, lines 55-59); and

- one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the requests to the interim fast-access-time cache to which the specified LAs are assigned (column 2, lines 11-15; column 4, lines 64-67).

Hicken et al teach,

- a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails (paragraph 0044).

Henry et al and Hicken et al are analogous art because they are from the same field of endeavor, namely memory storage systems.

At the time of invention, it would have been obvious to one of ordinary skill in the art, having both the teachings of Hicken et al and Henry et al to combine the redundant caches of Hicken et al into Henry et al for the benefit of restoring redundancy in the event of failure.

The Examiner notes that all of the claimed elements are known. The only difference is the combination of the known elements into one system. Implementing the redundant cache system of Hicken et al into Henry et al would have offered the system of Henry et al the benefit of having a mirror of the data within a cache should a failure occur. Upon such failure, redundancy could be reestablished (See Hicken et al Paragraph 0010-0011). Such combination would have yielded predictable results to

one of ordinary skill in the art at the time of invention thus obviating that which is instantly claimed.

Therefore it would have been obvious to combine Henry et al with Hicken et al to obtain the invention as specified in claims 70-71.

As per dependent claim 71, Henry et al. teach the storage system according to claim 70, wherein the mass storage devices comprise one or more disks (disks 1-5 of Fig. 2; column 5, lines 9-23).

As per independent claim 72, Henry et al. teach a method for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs) (Fig. 2; column 4, lines 10-23, lines 55-67), comprising:

- assigning to each of a plurality of interim-fast-access-time caches (cache pools 1 and 2), each of the plurality of interim fast-access-time caches comprising software (the Examiner notes that as the caches store data, the caches store software as instantly claimed) and configured to operate as controllers substantially independently of one another, a respective second range of the LAs (each cache pool is assigned separate LBA ranges for the associated disks; column 5, lines 45-55);
- coupling the plurality of interim-fast-access-time caches to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range (column 4, lines 55-59);
- receiving input/output (IO) requests from host processors directed to specified LAs (column 2, lines 11-15; column 4, lines 64-67); and

- directing all the IO requests to the interim-fast-access-time cache to which the specified LAs are assigned (column 2, lines 11-15; column 4, lines 64-67).

Hicken et al teach,

- wherein a further interim fast-access-time cache is adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data. from and provide data to the one or more mass storage devices having. LAs within the respective second range when any interim fast-access-time cache fails (paragraph 0044).

Henry et al and Hicken et al are analogous art because they are from the same field of endeavor, namely memory storage systems.

At the time of invention, it would have been obvious to one of ordinary skill in the art, having both the teachings of Hicken et al and Henry et al to combine the redundant caches of Hicken et al into Henry et al for the benefit of restoring redundancy in the event of failure.

The Examiner notes that all of the claimed elements are known. The only difference is the combination of the known elements into one system. Implementing the redundant cache system of Hicken et al into Henry et al would have offered the system of Henry et al the benefit of having a mirror of the data within a cache should a failure occur. Upon such failure, redundancy could be reestablished (See Hicken et al Paragraph 0010-0011). Such combination would have yielded predictable results to one of ordinary skill in the art at the time of invention thus obviating that which is instantly claimed.

Therefore it would have been obvious to combine Henry et al with Hicken et al to obtain the invention as specified in claims 72-73.

As per dependent claim 73, Henry et al. and Hicken et al teach the method according to claim 72, wherein the mass storage devices comprise one or more disks (disks 1-5 of Fig. 2; column 5, lines 9-23 of Henry et al).

Response to Arguments

Applicant's arguments, filed 4 May 2009, have been fully considered and are partially persuasive. However, upon further consideration, a new ground(s) of rejection is made as noted *supra*.

With respect to Applicant's argument located within the last paragraph of the second page of the instant remarks (numbered as page 15) and continuing through to the top of the third page of the instant remarks (numbered as page 16) which recites:

"In other words, detection of a failed storage controller is not the equivalent to detection a failed cache because to correct the situation (and eliminate the detection of a failure); the storage controller needs to be repaired or replaced, not the cache. Therefore, Hicken fails to disclose at least the elements of: "detecting an inability of the second cache to retrieve data from or store data at the first part of the first range of LAs," as recited in claim 41 because Hicken only discloses detecting the failure of a storage controller, not a cache."

The Examiner respectfully disagrees. While Applicant's comments are well taken and appreciated, the Examiner wishes to note the following. Hicken teach the failure of a storage controller. Nowhere does Hicken teach recovering the data from the cache of the failed storage controller. However, Hicken does teach that, upon failure of the storage controller, the mirrored cache that was mirroring the data in the cache of the failed storage controller is mirrored to a new cache (Paragraph 0011). This is done

such that the recovery of the redundancy of the data that was in the cache of the failed storage controller is reestablished. Thus, Hicken teach that which is instantly claimed. The Examiner further notes that in Paragraph 0012, Hicken explicitly teaches the failure of a cache and the recovery of the redundancy as noted above.

With respect to Applicant's argument located within the fourth paragraph of the third page of the instant remarks (numbered as page 16) which recites:

"Claims 26-33, which are now claims 66-73, respectively, stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,898,666 issued to Henry et al. ("Henry"). Applicants respectfully traverse the rejection because 1) the Patent Office has failed to establish a prima facie case of anticipation; and 2) Henry fails to disclose each and every element of claims 26-33."

The Examiner agrees and refers Applicants to the rejection made *supra* with respect to Henry et al in view of Hicken et al.

As the remained of Applicant's arguments are drawn to Hicken et al's alleged failure to teach the claimed limitation of, '*detecting an inability of the second cache to retrieve data from or store data at the first part of the first range of LAs*', the Examiner, in view of compact prosecution, incorporates by reference herein the comments made *supra* with respect to the first point of argument.

Any argument not specifically addressed is considered moot in view of the new ground(s) of rejection.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Matthew Bradley whose telephone number is (571) 272-8575. The examiner can normally be reached on 6:30-3:00 M-F.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Christian Chace can be reached on (571) 272-4190. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

CPC/mb
/Christian P. Chace/
Supervisory Patent Examiner, Art Unit 2187